

Leveraging Synthetic Targets for Machine Translation

Sarthak Mittal^{†1,2}

Oleksii Hrinchuk³

Oleksii Kuchaiev³

¹Mila, ²Universite de Montreal, ³NVIDIA

Abstract

In this work, we provide a recipe for training machine translation models in a limited resource setting by leveraging synthetic target data generated using a large pre-trained model. We show that consistently across different benchmarks in bilingual, multilingual, and speech translation setups, training models on synthetic targets outperforms training on the actual ground-truth data. This performance gap grows bigger with increasing limits on the amount of available resources in the form of the size of the dataset and the number of parameters in the model. We also provide preliminary analysis into whether this boost in performance is linked to ease of optimization or more deterministic nature of the predictions, and whether this paradigm leads to better out-of-distribution performance across different testing domains.

1 Introduction

Neural Machine Translation (NMT) (Bahdanau et al., 2014; Wu et al., 2016; Stahlberg, 2020) relies on deep learning models to train end-to-end translation systems. With the advent of deep recurrent models like LSTMs (Hochreiter and Schmidhuber, 1997; Sundermeyer et al., 2014; Chung et al., 2014) and their attention-augmented improvements (Bahdanau et al., 2014; Luong et al., 2015), these models outperformed traditional statistical (Koehn, 2009; Della Pietra, 1994) and rule-based (Lagarda et al., 2009; Nirenburg, 1989) approaches. Recently, with the introduction of fully attention-based networks (Vaswani et al., 2017; Dehghani et al., 2018; Sukhbaatar et al., 2019; Dai et al., 2019; Kitaev et al., 2020; Choromanski et al., 2020; Mittal et al., 2021) and increase in compute and data, large-scale Transformer networks have dominated the field of natural language process-

ing (Devlin et al., 2018; Brown et al., 2020; Hoffmann et al., 2022; Shoeybi et al., 2019), and machine translation in particular (Edunov et al., 2018; Raffel et al., 2020), leading to not only better performance but also more efficient training through their parallel computations (Ott et al., 2018).

While it has been established that scaling up data and compute boosts the performance of the large-scale NMT systems (Gordon et al., 2021; Ghorbani et al., 2021; Kaplan et al., 2020; Bahri et al., 2021), there is still a need to focus on budget models that can run on mobile and edge computing devices. In other tasks, like end-to-end speech translation, training data is scarce and expensive. Inspired from these needs, we provide a recipe for training text-to-text and text-to-speech translation systems in a limited resource setting at the modest overhead of running inference of pre-trained models on the source sentences.

Though in theory, increasing the amount of data provides a relatively simple methodology for bolstering the performance of current AI systems, it is difficult to do so when obtaining new data is costly, especially because of the labeling process in supervised learning. On the other hand, there have been a variety of approaches leveraging synthetic data to either improve the robustness of the systems or to boost their performance. This can be achieved by introducing adversarial examples in the training set (Goodfellow et al., 2014), considering knowledge distillation when provided access to large models but not their pre-training data (Buciluă et al., 2006; Gou et al., 2021; Hinton et al., 2015; Kim and Rush, 2016; Urner et al., 2011; Cheng et al., 2020; Phuong and Lampert, 2019; Tan et al., 2019), or using forward and back translation techniques (Zhang and Zong, 2016; Sennrich et al., 2015; Bogoychev and Sennrich, 2019; Edunov et al., 2018; Hoang et al., 2018) when additional monolingual data, which is easily available, is leveraged to generate synthetic targets to augment the amount of data.

[†] Work done during an internship at NVIDIA. Correspondence author sarthmit@gmail.com

We used the NeMo codebase (Kuchaiev et al., 2019) for all our experiments.

In this work, we use large, often pre-trained, NMT systems to provide synthetic targets which can be leveraged to train high performing low-compute and low-data models. Our findings show that using these synthetic translations to train different NMT systems leads to considerable improvements, even if we remove all ground-truth feedback from the training set. We also test models trained with synthetic targets on out-of-distribution settings through translations on different domains in a zero-shot manner as well as finetuning them on a different domain using additional synthetic data from an existing finetuned model, and further highlight the improvements obtained in both. We also find additional evidence to support that the ease of optimization resulting from training on synthetic targets does not completely explain its superior performance (He et al., 2019). Instead, we show case that models trained with synthetic targets are more deterministic than their counterparts which are trained on real targets.

Our key contributions are as follows

- We provide a recipe for training better models in resource and compute constrained settings provided access to a pre-trained system and validate it on bilingual and multilingual translation setups, as well as out-of-domain generalization.
- We provide analysis into the reasoning behind this improved performance and demonstrate that it is not solely because of ease of optimization but instead, we believe it is due to more deterministic nature of such systems.

2 Related Work

Synthetic targets have been consistently used in the past to either augment the amount of data available or to boost performance through knowledge transfer between models. In the domain of Machine Translation, a popular way of augmenting data is by considering monolingual data that is available in abundance and obtaining its paired counterpart using a trained model.

Back Translation. Back translation (Sennrich et al., 2015; Edunov et al., 2018; Hoang et al., 2018) relies on translating unseen sentences from the target side back to the source side using an existing trained model to provide new paired samples which can then be leveraged to train a model in the forward direction.. For example, if the task is to translate sentences from language \mathcal{S} to \mathcal{T} ($\mathcal{S} \rightarrow \mathcal{T}$),

one can obtain a corpus of monolingual data from \mathcal{T} and translate it backwards to \mathcal{S} using an existing trained $\mathcal{T} \rightarrow \mathcal{S}$ translation model. This would then provide additional paired data that can be combined with the existing ground-truth data to train a $\mathcal{S} \rightarrow \mathcal{T}$ translation model.

Forward Translation. Analogous to back translation, forward translation (Zhang and Zong, 2016; Bogoychev and Sennrich, 2019) or self-training (He et al., 2019) relies on training a standard translation model using which additional data is obtained by translating new sentences from the source to the target side and then re-training the translation model using this additional data. For example, to translate in the $\mathcal{S} \rightarrow \mathcal{T}$ direction, one first trains a model using existing data and then leverages this model to generate targets for a corpus of monolingual data from \mathcal{S} to provide new paired data. This data is then combined with the original data for re-training of the $\mathcal{S} \rightarrow \mathcal{T}$ translation model. Typically, forward translation is not as effective as back translation since the errors of the model are further propagated in the data in the former case (Burlot and Yvon, 2019).

Our approach can also be related to knowledge distillation, which has been the popular choice for transferring knowledge from a larger (called teacher) to a smaller model (called student) by enforcing similarities at different levels (Hinton et al., 2015; Gou et al., 2021; Freitag et al., 2017; Kim and Rush, 2016), eg. in the output or the representation space of the two models. We give a brief overview of the different strategies used.

Soft Target Matching. The earliest works on knowledge distillation transfer knowledge by enforcing the soft-logits of the student to be close to those of the teacher (Hinton et al., 2015). This is accomplished by introducing a loss term that penalizes deviation of the student model’s logits from the teacher, and this loss can either be used as is or added as a regularizing effect in training. This formulation of knowledge distillation has been revisited as *Word-Level Knowledge Distillation* for NMT in (Kim and Rush, 2016; Wang et al., 2021). While it is a simple way of distilling the teacher’s knowledge into the student, it can be computationally expensive if the number of classes, or equivalently the vocabulary \mathcal{V} , is large as it requires either storing all the $|\mathcal{V}|$ soft-logits for all words of the whole dataset or requires access to the teacher model on the fly, which can make training slow.

Model	Data	Real	Synthetic	Δ
2×2	0.5M	17.6	20.9	3.3
	5M	22.6	23.5	0.9
	25M	23.1	23.9	0.8
6×6	0.5M	18.0	22.5	4.5
	5M	24.7	25.7	1.0
	25M	25.3	25.9	0.6
24×6	0.5M	18.6	22.9	4.3
	5M	25.7	26.5	0.8
	25M	26.4	26.4	0.0

Table 1: English to Russian NMT models trained on subsets of WMT21 Machine Translation dataset. Synthetic targets are generated with 24×6 teacher model trained on the full 25M WMT dataset.

Representation Matching. (Romero et al., 2014; Zagoruyko and Komodakis, 2016; Lee et al., 2018; Heo et al., 2019; Passban et al., 2021; Chen et al., 2021) Another way of transferring knowledge from the teacher to the student is by matching their intermediate representations. Again, this can be accomplished by considering a regularization term $l_{reg}(g_t(\phi_t), g_s(\phi_s))$ where l_{reg} is some notion of similarity, ϕ_t, ϕ_s are the intermediate representations of the teacher and student respectively and g_t, g_s are functions that map the two representations to the same space, which is needed as the student is often smaller than the teacher model. While intuitively simple, this formulation is harder to implement and tune as the intermediate representations may be of different shapes, making it non-trivial to obtain a notion of similarity between the two. For example, if g_t and g_s map all the representations to the same point, the matching loss would be low even though the representations themselves ϕ_t, ϕ_s can be quite dis-similar.

Sequence-Level Knowledge Distillation. Kim and Rush (2016) propose *Sequence-Level Knowledge Distillation* which does not rely on soft-logits from the teacher model but instead relies on the synthetic translations obtained from the teacher model. Using synthetic targets is computationally efficient as the computation does not rely on matching the soft-logits across the whole of vocabulary but instead relies on sparse signals. Moreover, Kim and Rush (2016) showcase that using synthetic targets in LSTM-based systems lead to improved performance as opposed to the traditional knowledge distillation approach based on matching soft-logits.

While similar to forward translation and sequence-level knowledge distillation, our ap-

Model	Data	Real	Synthetic	Δ
2×2	0.5M	18.6	22.8	4.2
	5M	23.8	26.0	2.2
	57M	24.2	26.3	2.1
6×6	0.5M	18.7	23.9	5.2
	5M	25.1	27.8	2.7
	57M	26.6	28.4	1.8
24×6	0.5M	19.2	24.1	4.9
	5M	25.9	28.4	2.5
	57M	27.1	29.0	1.9

Table 2: English to German NMT models trained on subsets of WMT21 Machine Translation dataset. Synthetic targets are generated with 24×6 teacher model trained on large amounts of data in addition to the full 57M WMT training dataset.

proach differs by leveraging pre-trained translation models trained on large amounts of data for synthetic targets as opposed to training from scratch and then re-training. Further, we also consider setups where the amount of data used for the teacher and the student model is different, and where their model sizes can be similar.

3 Method

Our aim is to perform Machine Translation from a source language \mathcal{S} to a target language \mathcal{T} given some training data $\mathcal{D}_{\mathcal{S} \rightarrow \mathcal{T}} = \{(\mathbf{s}_i, \mathbf{t}_i)\}_{i=1}^N$ where $\mathbf{s}_i \in \mathcal{S}$ is the source sentence and $\mathbf{t}_i \in \mathcal{T}$ denotes the target sentence which is the ground-truth translation corresponding to \mathbf{s}_i . Further, we assume that we have access to a teacher network $f_{\mathcal{S} \rightarrow \mathcal{T}}(\cdot)$, using which we obtain synthetic targets $f_{\mathcal{S} \rightarrow \mathcal{T}}(\mathbf{s}_i)$ to construct the synthetic dataset $\mathcal{D}'_{\mathcal{S} \rightarrow \mathcal{T}} = \{(\mathbf{s}_i, f_{\mathcal{S} \rightarrow \mathcal{T}}(\mathbf{s}_i))\}_{i=1}^N$.

For our experiments, we consider different dataset sizes for the student models by subsampling from $\mathcal{D}_{\mathcal{S} \rightarrow \mathcal{T}}$ and $\mathcal{D}'_{\mathcal{S} \rightarrow \mathcal{T}}$ respectively. All the models considered in this work rely on the Encoder-Decoder Transformer architecture (Vaswani et al., 2017) with the teacher network generally having 24 encoder and 6 decoder layers (24×6). We consider different model sizes for the student, ranging from small models to matching the teacher’s size.

Typically, knowledge distillation considers the same input data for training both the teacher and the student. Instead, we perform analysis where the student has access to different amounts of data. Also, unlike knowledge distillation where knowledge is transferred from a bigger teacher to a smaller student network, we additionally consider setups

Model	Data	Real				Synthetic				Δ
		en-de	en-es	en-fr	Avg.	en-de	en-es	en-fr	Avg.	Avg.
2×2	1.5M	24.9	29.6	29.3	27.9	28.2	31.4	31.3	30.3	2.4
	15M	29.5	31.8	32.6	31.3	31.9	33.0	33.5	32.8	1.5
	300M	30.1	32.5	32.8	31.8	32.3	33.2	33.7	33.0	1.2
6×6	1.5M	25.9	30.1	30.2	28.7	31.0	32.5	32.5	32.0	3.3
	15M	33.2	33.9	35.0	34.0	36.0	34.4	35.4	35.2	1.2
	300M	34.5	34.4	35.6	34.8	36.1	34.5	35.8	35.4	0.6
24×6	1.5M	25.8	30.0	29.6	28.5	31.8	33.3	33.3	32.8	4.3
	15M	35.0	34.5	35.7	35.0	36.9	34.9	36.4	36.0	1.0
	300M	36.1	34.9	36.8	35.9	37.9	35.2	37.0	36.7	0.8

Table 3: Multilingual Machine Translation systems trained to translate sentences from English to three different languages: German, Spanish and French. Synthetic targets are generated with 24×6 bilingual teachers trained on a large corpus of parallel data. SacreBLEU on WMT20 dev set is reported.

where the student network is as big as the teacher, and showcase the benefits in this regime.

4 Experiments

For all our experiments, we only considered Transformer models for both teachers and students. Unless specified, we used the Pre-Layer Normalization variant where LayerNorm is applied before the respective attention and residual computations (Xiong et al., 2020). In our experiments, we consider two text-to-text machine translation setups: bilingual and multilingual, and one speech-to-text setup. For the text-to-text machine translation experiments, we considered byte-pair encoding (Britz et al., 2017) for bilingual experiments and sentence-piece encoding (Kudo and Richardson, 2018) for multilingual experiments.

4.1 Bilingual Machine Translation

We first test the benefits of training with synthetic targets on bilingual machine translation where models are trained to translate sentences from one specified (source) language to another (target) language. We conducted experiments with the source language as English and the target languages as Russian and German. We consistently see improvements when training with synthetic targets, and these benefits are substantial when the student is trained on limited data. Even on the same amount of data, we see benefits of using synthetic data when the student has lower complexity.

English to Russian. We used the WMT’21 dataset for English to Russian Machine Translation, where we trained the models with different tokenizers and vocabularies on the source and target side with the vocabulary size of 24576. For the

teacher model, we trained a baseline Transformer with 24 encoder and 6 decoder layers. For the student, we consider two different axis of analysis: models with lower capacity than the 24×6 teacher, and models trained with fewer data subsampled from the training set used to train the teacher.

In Table 1, we highlight that in the low-data regime for the student, consistently across all the different model sizes, training solely on synthetic targets obtained from the teacher model leads to much better performance than training on real ground-truth data. We also see that even when keeping the amount of data fixed (25M sentences; which is the same data on which the teacher model was trained), we see improvements on using synthetic targets in training smaller models. Thus, the only avenue where we don’t get a substantial improvement is when the student uses the same dataset as the teacher model and has high complexity, similar to the teacher. This, however, is intuitive and doesn’t pose a problem since our aim is for better low-compute models.

English to German. We consider the WMT’21 dataset for the task of English to German Machine Translation, where we trained the models with shared tokenizers and vocabularies on the source and target side, with the vocabulary size of 32000. For the teacher model, we picked a published transformer model with 24 encoder and 6 decoder layers from the NeMo codebase (Kuchaiev et al., 2019) which was trained on substantially more data than WMT. For the student models, we again consider two different axis of analysis: models with lower capacity than the 24×6 teacher, and models trained with different percentages of the WMT data.

In Table 2, we see that consistent with our En-

Model	Data	Real				Synthetic				Δ
		de-en	es-en	fr-en	Avg.	de-en	es-en	fr-en	Avg.	Avg.
2×2	1.5M	28.4	29.1	32.2	29.9	31.9	30.8	34.1	32.3	2.4
	15M	34.6	32.8	36.2	34.5	37.1	33.6	37.5	36.1	1.6
	300M	35.5	32.7	36.9	35.0	38.3	33.8	37.6	36.6	1.6
6×6	1.5M	30.8	30.1	33.2	31.4	35.3	31.8	35.3	34.1	2.7
	15M	38.1	34.0	38.3	36.8	40.2	35.0	39.1	38.1	1.3
	300M	40.0	34.2	38.8	37.7	41.3	35.5	39.6	38.8	1.1
24×6	1.5M	31.4	30.5	33.8	31.9	36.0	32.8	36.2	35.0	3.1
	15M	38.8	34.6	39.6	37.6	40.8	35.6	40.1	38.8	1.2
	300M	40.1	35.6	40.9	38.9	41.8	36.3	41.2	39.8	0.9

Table 4: Multilingual Machine Translation systems trained to translate sentences from German, Spanish, and French to English. Synthetic targets are generated with 24×6 bilingual teachers trained on a large corpus of parallel data. SacreBLEU on WMT20 dev set is reported.

English to Russian experiments, in low data regime the student models trained with synthetic targets outperform the models trained on real ground-truth data. In particular, since the teacher model was trained on an enormous corpus outside of WMT as well, we see that even on using a 24×6 transformer model which has the same complexity as the teacher on the full WMT21 dataset (57M sentences), it is still beneficial to train it with synthetic data as opposed to real data. Additionally we can also see that smaller models (6×6) trained on synthetic targets outperform larger models (24×6) trained on real targets, while also providing faster training and inference.

Our English to Russian and English to German experiments show that indeed across different amounts of data regime and model complexities, as long as one does not approach the large-scale teacher model in both the regimes, it is beneficial to train systems using synthetic targets from the pre-trained large-scale model as opposed to training it on the ground-truth targets. We show that through this simple recipe, one can train much better smaller scale models which is beneficial to have when considering deployment in low-resource settings where inference latency needs to be low. It also provides a viable strategy for training of models when access to large-scale systems is provided but their training data is not.

4.2 Multilingual Machine Translation

Next, we move our attention towards multilingual machine translation where a single model is trained to perform translation from multiple different source languages to various different target languages. In particular, we focus on two different

multilingual settings; translating sentences from (a) English to German, Spanish and French, and (b) German, Spanish and French to English. An important difference from bilingual experiments is that in this case, we obtain synthetic targets for the student multilingual models using published bilingual teachers. As in the bilingual setup, we again see consistent improvements of using synthetic targets instead of real ground-truth targets.

English to German/Spanish/French. We consider the setup of training a single model to translate English sentences into three different languages: German, Spanish and French. These models are trained with shared tokenizers and vocabularies on the source and target side, with the vocabulary size of 64000. We query bilingual teachers trained on considerably large datasets to obtain synthetic targets for training the multilingual translation systems and compare it to training done on real ground-truth translations. We consider three different dataset sizes for the multilingual setup with 0.5M, 5M, and 100M sentences for each language.

In Table 3, we see that across different dataset sizes (1.5M = 0.5M for each pair; similar for 15M and 300M) and model complexities, training on synthetic targets outperforms training on ground-truth data. Thus, in the presence of large pre-trained bilingual experts, this provides a recipe for training stronger and more powerful multilingual models of various sizes.

German/French/Spanish to English. We perform similar analysis in the reverse direction where we train multilingual models to translate sentences from German, French and Spanish to English. We use shared tokenizers and vocabularies, with the vocabulary size of 64000, and use bilingual experts

German targets source	Size	Must-C v2		IWSLT tst							
		dev	tst	2010	2013	2014	2015	2018	2019	2020	Avg.
Real	590K	25.2	27.8	22.1	27.9	23.8	21.3	20.9	20.1	20.9	22.4
Synthetic, WMT21 teacher	590K	28.3	28.9	24.5	28.0	23.9	23.0	22.6	21.7	23.3	23.9
+ fine-tuned on IWSLT	590K	29.2	30.2	25.0	29.5	24.8	24.9	24.2	23.4	25.3	25.3
+ extra ASR data	1.25M	30.6	31.0	27.2	31.3	27.4	25.8	25.1	24.3	26.4	26.8

Table 5: Speech Translation systems trained to translate English audio to German text. Synthetic targets are generated with 24×6 bilingual teacher trained on WMT21 and fine-tuned on IWSLT22.

trained on large datasets to provide synthetic targets. We train the multilingual models on synthetic and ground-truth targets with 0.5M, 5M, and 100M sentences for each language pair.

We highlight the results in Table 4 for different dataset and model sizes and see consistent improvements on using synthetic targets as opposed to real ground-truth targets for training.

Our experiments on multilingual translation in both directions reveal that when we use bilingual experts for each language pair as teachers, we see consistent improvements when training solely on synthetic data. We believe that this can pave the way for utilization of bilingual models for better and more efficient training of multilingual models, given bilingual models currently outperform multilingual models while the latter are more memory efficient as only a single model needs to be stored.

4.3 Speech Translation

In end-to-end speech translation (ST), the task is to train a model which translates speech in one language into a text in another. In contrast to text-to-text NMT, data for this particular task is much more scarce and expensive. However, this problem can be solved with the help of readily available text-to-text models and a large corpora of ASR data. Surprisingly, completely replacing all target language transcripts with synthetic data improves the performance of the ST models we trained.

Our speech translation models consist of a 17-layer Conformer encoder (Gulati et al., 2020) initialized with pre-trained speech recognition (ASR) encoder followed by a 6-layer Transformer decoder initialized randomly. For training, we used all available En→De ST datasets from IWSLT’22 competition (Anastasopoulos et al., 2022) which amounted to 590K examples after cleaning. To generate synthetic targets, we used 24×6 teacher model trained on WMT21 with the optional in-domain fine-tuning on 250K sentences from Must-C v2 dataset (Cattoni et al., 2021).

Model	Data	Real	Synthetic	Δ
2×2	IWSLT	27.9	30.4	2.5
	Medical	27.3	30.6	3.3
	Law	36.5	39.0	2.5
6×6	IWSLT	30.0	31.6	1.6
	Medical	30.1	31.7	1.6
	Law	39.8	40.8	1.0
24×6	IWSLT	31.2	33.1	1.9
	Medical	31.3	33.1	1.8
	Law	40.4	42.1	1.7

Table 6: English to German NMT model trained on WMT21 Machine Translation dataset with real and synthetic targets respectively, evaluated out-of-distribution on different domains. The teacher model used to generate synthetic targets was trained on additionally large amounts of data.

In Table 5, we see that replacing real targets with synthetic ones leads to over 1.5 BLEU score improvement, even though the teacher model was trained on out-of-domain data. When teacher model is fine-tuned in-domain, the score improvement goes to 3 BLEU. Finally, using the synthetic translations, we can expand our dataset by translating ASR-only datasets from IWSLT’22 which do not have German translations. Adding additional 660K examples leads to another 1.5 BLEU improvement over, even though the additional data was out of domain to TED talks we evaluate on.

4.4 Out-of-Domain Evaluation

One might argue that the model trained on a particular dataset overfits to it and using it to translate sentences from other domains will produce poor results, which can be exacerbated on using synthetic targets from models trained on a particular domain. In the next series of experiments, we evaluate out-of-domain performance of models trained with synthetic data.

Table 6 shows the performance of models trained on WMT’21 dataset evaluated on three different non-news domains: TED talks (IWSLT), medical,

Fine-tuning targets	2×2	6×6	24×6
<i>Training from scratch</i>			
Real	16.9	17.0	19.2
Synthetic, WMT Teacher	20.6	21.9	22.7
+ fine-tuned on IWSLT	21.8	22.9	25.9
<i>Pre-training on real data</i>			
Real	30.9	33.6	34.5
Synthetic, WMT Teacher	30.9	32.6	33.7
+ fine-tuned on IWSLT	32.5	34.1	35.5
<i>Pre-training on synthetic data</i>			
Real	32.4	33.8	34.9
Synthetic, WMT Teacher	32.0	33.3	33.6
+ fine-tuned on IWSLT	33.4	35.1	35.6

Table 7: English to German NMT model trained on WMT21 dataset and additionally fine-tuned on IWSLT dataset. SacreBLEU on Must-C t_{st} -COMMON dataset is reported.

and law. As we see, training on synthetic targets is beneficial here as well even though all training was done on the WMT dataset.

4.5 In-Domain Fine-tuning

Finally, we utilize synthetic targets for in-domain fine-tuning, where pre-training is done on real or synthetic data and then further fine-tuning is done using real, out-of-domain synthetic and in-domain synthetic data. In Table 7, we train NMT models on WMT’21 dataset with either real or synthetic targets and then additionally fine-tune them on IWSLT data which is in-domain to evaluation t_{st} -COMMON dataset comprising of TED talks.

We see that fine-tuning on synthetic targets generated with out-of-domain model actually hurts the model performance but pre-training on it works well. Also, fine-tuning on synthetic targets generated with in-domain model is superior to fine-tuning on real in-domain data no matter what data was used for pre-training. Training from scratch with synthetic targets generated by both models outperforms real targets by a large margin.

5 Ablations

Our key ablations into understanding the phenomena presented involve analysing whether the results can be solely explained from the lens of optimization, or are there other reasons at play (eg. stochasticity of the predictive model).

5.1 Optimization Problem

It can be argued that training on teacher outputs is easier (He et al., 2019) and its superior performance




Figure 1: Training loss of models trained with synthetic and real targets on the English to German machine translation. Both models are 6×6 Transformers; synthetic targets are obtained with a 24×6 teacher model.

is just an artifact of a well-behaved optimization landscape leading to “better solution spaces”. Figure 1 shows that synthetic targets are indeed easier to fit as both the training loss and its variance are lower. However, we do not agree that training on synthetic targets leads to better local optima on which real-data training can capitalize on.

To see this, we pre-train some of the systems from Section 4 on synthetic targets and half-way during the training, switch out to real ground-truth targets. Our hypothesis is that if training on synthetic targets leads to better solution spaces, using the corresponding model parameters as initialization will not hurt the overall performance of training with real data, as evaluation data is considered to be from the same distribution as the real data and not the synthetic one.

However, our experiments in Figure 2 imply that when we switch the data from synthetic to real, we do see an immediate drop in performance. Moreover, we also see that when we train on real targets and switch to synthetic targets, we get an immediate boost in performance. Together, the two results indicate that training with synthetic data does not lead to better solutions to the underlying optimization problem and that the reasoning behind its workings is something else.

5.2 Top-k Performance

Next, we analyze whether the models trained with synthetic targets are more deterministic than those trained with real data. Our hypothesis is that training on synthetic targets dampens some of the noise present in the ground-truth dataset. To explore this claim, we use the models trained in Section 4 and evaluate them at different levels of top-k sampling for inference and compare it with models trained on ground-truth data.




Figure 2: We highlight that the benefits of training using a teacher model are not because of optimization. In particular, we train models on synthetic targets and then mid-way between training, we switch to real targets. We see that this leads to a degradation in performance implying that synthetic training does not lead to better initialization for the model, and this is consistent over different setups. Further, if we train on real targets and switch to synthetic ones, we see an immediate boost in performance.

Language	Targets	Predictive Entropy		
		2×2	6×6	24×6
English to German	Real	2.4	2.4	2.5
	Synthetic	1.8	1.8	1.7
English to Russian	Real	2.6	2.4	2.3
	Synthetic	1.8	1.8	1.9

Table 8: Average entropy of the predictive distribution for translation $p(y|x)$ is lower for models trained with synthetic targets than those trained with real targets, implying more deterministic nature of the former.

Our findings in Figure 3 showcase that the drop in performance with increasing k in top- k is less in models trained with synthetic targets than in those trained with real targets. This does highlight that models trained with synthetic targets capture less noise as the degradation of performance is less when we make the sampling more noisy.

5.3 Predictive Entropy

As a final analysis of whether models trained with synthetic targets are more deterministic, we compute the predictive entropy of the distribution over the logits that predict the next token in translation. Our findings in Table 8 highlight that indeed the predictive entropy of models trained with synthetic targets is lower than the models trained with ground-truth targets, implying more deterministic nature of the former translation systems.

Together, our analysis on top- k performance and predictive entropy provide some evidence that the models trained with synthetic targets are more deterministic and hence are more robust and perform better, even on out of distribution shifts.

6 Conclusion

Inspired by the recent advances in knowledge distillation and the need for better performing low-




Figure 3: Degradation of Machine Translation performance with increasing k in top- k sampling for the decoding procedure for both English to German and English to Russian translations using student 24×6 models trained with real and synthetic targets respectively.

resource and low-compute models, we provide a recipe that leverages large-scale pre-trained translation systems as teacher models which provide synthetic targets for training of smaller and low-resource models. Surprisingly, we see a considerable increase in the performance of smaller models when only teacher outputs are provided as opposed to any proportion of real ground-truth translations. We also see additional benefits of using synthetic targets for training, namely faster convergence and improved translations with top- k sampling when compared to models trained solely on real ground-truth translations.

This improvement in performance of small or low-resource models comes at the additional inference costs tied with the large teacher model. However, this is a one-time data curating cost based on which training of multiple different smaller models can be accomplished which will enjoy not only faster inference (owing to their smaller size) but also better performance than if they were trained on the original data, sometimes even better than larger models trained on real data. We believe this is an exciting avenue of research as low-compute high-performance models are important when deployed in constrained settings like edge computing and mobile devices.

Ethics Statement

We provide a methodology for improving the performance of resource and data constrained translation systems which rely on obtaining synthetic targets from larger pre-trained systems. Given the dependence on large pre-trained systems, we believe that their biases can negatively impact the biases and fairness of the smaller consecutively trained systems. This is a problem common with any type of knowledge transfer where the biases of the base model can also be transferred to the student system and approaches on mitigating biases in the larger models in the first place would be a potential solution that can alleviate this problem.

References

- Antonios Anastasopoulos, Loïc Barrault, Luisa Bentivogli, Marceley Zanon Boito, Ondřej Bojar, Roldano Cattoni, Anna Currey, Georgiana Dinu, Kevin Duh, Maha Elbayad, et al. 2022. Findings of the iwslt 2022 evaluation campaign. In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 98–157.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. 2021. Explaining neural scaling laws. *arXiv preprint arXiv:2102.06701*.
- Nikolay Bogoychev and Rico Sennrich. 2019. Domain, translationese and noise in synthetic data for neural machine translation. *arXiv preprint arXiv:1911.03362*.
- Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. 2017. Massive exploration of neural machine translation architectures. *arXiv preprint arXiv:1703.03906*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Cristian Bucilua, Rich Caruana, and Alexandru Niculescu-Mizil. 2006. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541.
- Franck Burlot and François Yvon. 2019. Using monolingual data in neural machine translation: a systematic study. *arXiv preprint arXiv:1903.11437*.
- Roldano Cattoni, Mattia Antonino Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. Mustc: A multilingual corpus for end-to-end speech translation. *Computer Speech & Language*, 66:101155.
- Defang Chen, Jian-Ping Mei, Yuan Zhang, Can Wang, Zhe Wang, Yan Feng, and Chun Chen. 2021. Cross-layer distillation with semantic calibration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7028–7036.
- Xu Cheng, Zhefan Rao, Yilan Chen, and Quanshi Zhang. 2020. Explaining knowledge distillation by quantifying the knowledge. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12925–12935.
- Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Kingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. 2020. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2018. Universal transformers. *arXiv preprint arXiv:1807.03819*.
- Vincent J Della Pietra. 1994. The mathematics of statistical machine translation: Parameter estimation. *Using Large Corpora*, page 223.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. *arXiv preprint arXiv:1808.09381*.
- Markus Freitag, Yaser Al-Onaizan, and Baskaran Sankaran. 2017. Ensemble distillation for neural machine translation. *arXiv preprint arXiv:1702.01802*.
- Behrooz Ghorbani, Orhan Firat, Markus Freitag, Ankur Bapna, Maxim Krikun, Xavier Garcia, Ciprian Chelba, and Colin Cherry. 2021. Scaling laws for neural machine translation. *arXiv preprint arXiv:2109.07740*.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Mitchell A Gordon, Kevin Duh, and Jared Kaplan. 2021. Data and parameter scaling laws for neural machine translation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5915–5922.
- Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. 2021. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Junxian He, Jiatao Gu, Jiajun Shen, and Marc’Aurelio Ranzato. 2019. Revisiting self-training for neural sequence generation. *arXiv preprint arXiv:1909.13788*.

- Byeongho Heo, Minsik Lee, Sangdoon Yun, and Jin Young Choi. 2019. Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3779–3787.
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7).
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. Iterative back-translation for neural machine translation. In *Proceedings of the 2nd workshop on neural machine translation and generation*, pages 18–24.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*.
- Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*.
- Philipp Koehn. 2009. *Statistical machine translation*. Cambridge University Press.
- Oleksii Kuchaiev, Jason Li, Huyen Nguyen, Oleksii Hrinchuk, Ryan Leary, Boris Ginsburg, Samuel Kriman, Stanislav Beliaev, Vitaly Lavrukhin, Jack Cook, et al. 2019. Nemo: a toolkit for building ai applications using neural modules. *arXiv preprint arXiv:1909.09577*.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Antonio-L Lagarda, Vicent Alabau, Francisco Casacuberta, Roberto Silva, and Enrique Diaz-de Liano. 2009. Statistical post-editing of a rule-based machine translation system. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*, pages 217–220.
- Seung Hyun Lee, Dae Ha Kim, and Byung Cheol Song. 2018. Self-supervised knowledge distillation using singular value decomposition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 335–350.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Sarthak Mittal, Sharath Chandra Rapparthi, Irina Rish, Yoshua Bengio, and Guillaume Lajoie. 2021. Compositional attention: Disentangling search and retrieval. *arXiv preprint arXiv:2110.09419*.
- Sergei Nirenburg. 1989. Knowledge-based machine translation. *Machine Translation*, 4(1):5–24.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. *Scaling neural machine translation*. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Peyman Passban, Yimeng Wu, Mehdi Rezagholizadeh, and Qun Liu. 2021. Alp-kd: Attention-based layer projection for knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13657–13665.
- Mary Phuong and Christoph Lampert. 2019. Towards understanding knowledge distillation. In *International Conference on Machine Learning*, pages 5142–5151. PMLR.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*.
- Felix Stahlberg. 2020. Neural machine translation: A review. *Journal of Artificial Intelligence Research*, 69:343–418.
- Sainbayar Sukhbaatar, Edouard Grave, Piotr Bojanowski, and Armand Joulin. 2019. Adaptive attention span in transformers. *arXiv preprint arXiv:1905.07799*.

- Martin Sundermeyer, Tamer Alkhouli, Joern Wuebker, and Hermann Ney. 2014. Translation modeling with bidirectional recurrent neural networks. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 14–25.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461*.
- Ruth Urner, Shai Shalev-Shwartz, and Shai Ben-David. 2011. Access to unlabeled data can speed up prediction time. In *ICML*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Fusheng Wang, Jianhao Yan, Fandong Meng, and Jie Zhou. 2021. Selective knowledge distillation for neural machine translation. *arXiv preprint arXiv:2105.12967*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. 2020. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*, pages 10524–10533. PMLR.
- Sergey Zagoruyko and Nikos Komodakis. 2016. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*.
- Jiajun Zhang and Chengqing Zong. 2016. Exploiting source-side monolingual data in neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545.

Appendix

A Implementation Details

For all our experiments, we rely on the NeMo codebase published in [Kuchaiev et al. \(2019\)](#). We do not perform extensive hyperparameter selection and instead just rely on the defaults provided. All the models that we train from scratch use the pre layer-norm transformer variant ([Xiong et al., 2020](#)) and are trained with 0.001 learning rate with a linear warmup followed by an exponential decay. For all the synthetic targets, we use beam search with beam size 4 for generating translations. All experiments also use label smoothing of 0.1 with a dropout of 0.1 as well. We only vary the models in their depth, while keeping the attention layer of 512 dimensions, feed-forward residual connections of 2048 and 8 attention heads, as is typical with transformer models. All experiments were done on 16 GPUs and for 150,000 iterations by when convergence of all models had been achieved. All the results are reported by considering single runs of the model as our experiments revealed very low variance between different runs.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
In Section 4.1, where we show that our approach doesn't provide benefits on improving models that have the same data and compute complexity as the teacher.
- A2. Did you discuss any potential risks of your work?
Not applicable. Our approach is more of an up-to-date analysis into some of the existing techniques used, and does not pose any risks beyond the risks of general improvement of AI.
- A3. Do the abstract and introduction summarize the paper's main claims?
Please refer to Abstract and Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Not applicable. Left blank.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Yes, we used the standard splits in WMT data and mentioned them in the paper.

C Did you run computational experiments?

Section 4; Experiments.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
We provide details about the framework used, as well as the size of models in terms of number of layers. Details in Appendix A

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

We did not do any extensive hyperparameter search and used default parameters. Details in Appendix A

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

We report a single run. Details in Appendix A.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 4 and first page.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Not applicable. Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Not applicable. Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Not applicable. Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.